

HillSense AI: Smart IoT-Based Soil Monitoring and Automated Irrigation System

Anmol Ben Emmanuel
Dept. of Computer Science and
Engineering
Tula's Institute, Dehradun

Praveen Negi
Dept. of Computer Science and
Engineering
Tula's Institute, Dehradun

Prajwal Panwar
Dept. of Computer Science and
Engineering
Tula's Institute, Dehradun

Yuvraj Singh Pokhariya
Dept. of Computer Science and
Engineering
Tula's Institute, Dehradun

Suraj Singh
Dept. of Computer Science and
Engineering
Tula's Institute, Dehradun

Dr. Saurabh Singh
Dept. of Computer Science and
Engineering
Tula's Institute, Dehradun

Abstract - This paper presents HillSense AI, a smart IoT-based soil monitoring and automated irrigation system built using an ESP32 microcontroller, multiple environmental sensors, and a cloud-connected AI recommendation engine. The system continuously reads soil moisture, temperature, humidity, and ambient light values, pushes them to Firebase Realtime Database over Wi-Fi, and makes them accessible through both a Flutter mobile application and a Streamlit web dashboard.

A TensorFlow/Keras neural network trained on over 100,000 agricultural records classifies soil quality into three categories — Poor, Medium, and Good — with a final training accuracy of 82.14% and a test accuracy of 93.42%.

Beyond classification, the system generates context-aware irrigation, crop suitability, and soil nutrient recommendations in real time. Automated irrigation is handled by a MOSFET-controlled DC water pump that switches on whenever moisture drops below 30% during daylight and shuts off after a fixed watering window to prevent overwatering. Across 30 independent test runs, the irrigation logic achieved a 100% activation success rate, and controlled watering reduced estimated water consumption by roughly 22–25% compared with fixed-schedule methods. A Google Cloud Run backend hosts the inference server, while Firebase timestamps provide continuous device-health monitoring so that disconnections are flagged immediately on every client. The result is an affordable, end-to-end smart agriculture platform that unifies sensor hardware, cloud services, machine learning, and user-facing interfaces in a single deployable solution.

Keywords - Smart Agriculture, Internet of Things, ESP32, Automated Irrigation, Soil Monitoring, CNN, TensorFlow, Firebase, Flutter, Precision Farming.

I. INTRODUCTION

Agriculture has always been central to human survival, but the gap between what modern technology can do and what actually reaches the average farmer keeps widening. Water scarcity, unpredictable monsoons, and soil degradation are problems that show up every season, yet most smallholder farmers still rely on intuition and fixed watering schedules rather than data. Over-irrigation washes away fertilizer and damages root zones; under-irrigation stunts growth and cuts yield. Neither outcome is acceptable when inputs are expensive and margins are thin.

The Internet of Things has made it genuinely practical to put real-time sensors in the field and act on their readings automatically. Microcontrollers like the ESP32 are powerful enough to handle Wi-Fi communication, sensor interfacing, and even lightweight decision logic at a price point that does not require institutional funding. When you pair that hardware with a cloud database and a simple mobile interface, farmers can see what is happening in their soil from wherever they happen to be — and the system can take corrective action without waiting for human input.

HillSense AI was built with that practical gap in mind. The core idea is straightforward: measure the things that actually determine whether a crop thrives — soil moisture, temperature, humidity, and light — and use those readings to drive two outputs: automated irrigation control and AI-generated recommendations about irrigation timing, crop selection, and soil nutrition. Everything runs through Firebase so the data is live on both a web dashboard and a Flutter mobile app, and a TensorFlow model deployed on Google Cloud Run handles the heavier classification work that goes beyond simple threshold logic.

A. Problem Statement

Traditional irrigation depends on human observation, which is inherently intermittent. A farmer who checks the field once in the morning may not catch the point in the afternoon when moisture crosses a critical threshold. Fixed schedules compensate for this by watering on a timer, but they ignore what the soil is actually doing — watering even when it has rained, skipping on a hot dry day because the schedule says so. The downstream effects are real: wasted water, leached nutrients, root disease from waterlogging, or stress from drought. What is missing is a system that monitors continuously, responds immediately, and gives the farmer enough information to make better decisions even when they are not physically present.

B. Objectives

The project set out to build a low-cost, integrated smart farming platform that monitors soil and environmental conditions in real time, classifies soil health using a trained neural network, generates practical agricultural recommendations, and automates irrigation without requiring manual intervention. A secondary objective was to make everything accessible remotely through a mobile application, so the farmer does not need to be at the field to know what is happening or to verify that irrigation has occurred.

II. LITERATURE REVIEW

Most early IoT irrigation work focused on the hardware side: put a moisture sensor on an Arduino, set a threshold, drive a relay. Muthuramalingam et al. showed that this basic pattern could cut manual watering effort significantly and improve water utilization, but the systems they described stopped at irrigation control and offered no analysis of what the underlying soil condition actually was [1].

The push toward AI-assisted agriculture picked up seriously around 2018. Liakos et al. surveyed machine learning use across the agricultural value chain and found strong results in soil classification, yield forecasting, and disease detection, though most published systems were tested under controlled conditions rather than real field deployments [2]. Kamilaris and Prenafeta-Boldú looked specifically at deep learning and noted that convolutional architectures, originally developed for images, adapt surprisingly well to tabular environmental data when the feature set is rich enough [3]. Wolfert et al. framed the broader challenge as a data pipeline problem: sensors generate useful readings, but the value only materializes when those readings are stored, processed, and presented in a form farmers can act on [4].

Cloud connectivity changed what was achievable at low cost. Firebase in particular became a practical choice for student and small-team projects because it handles real-time synchronization across web and mobile clients without requiring a dedicated server [5]. Elijah et al. reviewed IoT architectures for agriculture and found that cloud-backed systems consistently outperformed local-only setups in terms

of accessibility and data persistence, though they noted that connectivity assumptions can be problematic in rural areas [6]. Khanna and Kaur argued that meaningful precision agriculture adoption would require solutions affordable enough for smallholder farmers, not just large commercial operations [7].

The gap that HillSense AI attempts to address is the integration gap. Existing systems tend to do one thing well — monitor, or irrigate, or recommend — but rarely bring all three together in a single deployable package with both a mobile interface and a trained model producing soil-quality predictions rather than just threshold alerts.

III. METHODOLOGY

The system architecture divides naturally into five layers that each handle a distinct part of the pipeline: sensor data collection, ESP32-side processing and irrigation control, cloud synchronization via Firebase, AI inference on Google Cloud Run, and user interfaces on web and mobile. These layers are described in detail below.

A. Sensor Data Collection

Three sensors connect to the ESP32. A capacitive soil moisture sensor measures volumetric water content as a percentage; capacitive sensing was chosen over resistive probes because it does not corrode in wet soil and gives stable long-term readings. A DHT11 module measures both ambient temperature in degrees Celsius and relative humidity as a percentage. An LDR provides a simple binary daylight/night signal that the irrigation logic uses to prevent nighttime watering, which tends to promote fungal growth and is generally not needed because evaporation is low. The ESP32 reads all three sensors on a continuous polling loop, averages multiple analog samples to reduce noise, and converts raw ADC values to calibrated units before doing anything else with them.

B. ESP32 Processing and Irrigation Control

Once readings are in calibrated form, the ESP32 applies the irrigation decision logic locally without waiting for a cloud round-trip. The rule is simple on paper but was tuned carefully during testing: if soil moisture is below 30% and the LDR reports daylight, the MOSFET gate is driven high, turning the pump on. The pump runs for a fixed window of five seconds, then the gate drops and the firmware waits before re-evaluating. This prevents the pump from running continuously even if moisture has not recovered yet, which would cause waterlogging. After the wait period, the cycle repeats. If either condition fails — moisture is adequate, or it is night — the pump stays off. The MOSFET driver was chosen over a relay specifically because relay-controlled pumps showed startup trigger instability in early testing; the MOSFET switches cleanly at ESP32 voltage levels and has no mechanical contact to wear out.

All sensor readings, pump state, light status, and a timestamp are then packaged and pushed to Firebase over the ESP32's built-in Wi-Fi. The Firebase library handles reconnection automatically, so brief signal drops do not lose data permanently.

C. Firebase Cloud Synchronization

Firebase Realtime Database acts as the central data bus for the whole system. Every time the ESP32 writes a new reading, Firebase propagates it to all connected clients in under a second. The database structure keeps things flat: a sensor node holds the current temperature, humidity, moisture, light status, pump status, and a last seen Unix timestamp. A separate predictions node holds the most recent output from the AI backend, and a recommendations node stores the three recommendation strings — irrigation, crop, and soil nutrient — that get refreshed with each inference cycle. Device health monitoring is implemented by comparing the current server time against last seen; if the gap exceeds a configurable threshold, the dashboard marks the device as OFFLINE and surfaces that status visibly so the farmer knows something needs attention.

D. Machine Learning Model

The soil quality classification model is a fully connected neural network built with TensorFlow and Keras. It takes seven input features: temperature, humidity, soil moisture, light intensity, pH level, nitrogen content, and potassium content. The first dense layer has 256 units with ReLU activation, followed by a dropout layer at rate 0.3 to reduce overfitting. A second dense layer with 128 units and the same dropout pattern feeds into a 64-unit layer, and finally a 3-unit softmax output layer that produces probabilities for the three soil quality classes: Poor, Medium, and Good. The model has approximately 43,000 trainable parameters in total.

Training used a dataset of over 100,000 agricultural records sourced from publicly available agricultural databases and augmented with synthetic samples to balance class distribution. An 80/20 train-test split was applied after label encoding and standard scaling of all input features. The Adam optimizer with a learning rate of 0.001 minimized sparse categorical cross-entropy over 50 epochs with a batch size of 32 and a patience-10 early stopping callback. Final training accuracy settled at 82.14% and test accuracy reached 93.42%, which is a meaningful gap that reflects generalization rather than memorization.

The trained model, scaler, and label encoder were saved as artifacts (`soil_quality.keras`, `scaler.pkl`, `label_encoder.pkl`) and packaged in a Flask-based REST API deployed on Google Cloud Run. When the ESP32 writes new sensor data to Firebase, the Cloud Run service reads the values, runs inference, writes the predicted soil quality class back to the prediction's node, and simultaneously generates the three recommendation strings based on the class and the raw sensor

context. The entire inference round-trip typically completes in under two seconds.

E. User Interfaces

Two interfaces expose the system to end users. The Streamlit web dashboard pulls from Firebase every few seconds using an auto-refresh loop and renders live sensor cards, Plotly gauge charts for temperature, humidity, and moisture, a soil status banner, three AI recommendation panels, an irrigation monitoring section that shows pump state, and historical trend charts. The Flutter mobile application does the same through Firebase's Flutter SDK, presenting a clean card-based layout with live sensor values, system status indicators, and the same three recommendation categories. Both interfaces surface the device ONLINE/OFFLINE status prominently so a disconnected ESP32 is never silently missed.

IV. TECHNOLOGIES USED

The hardware side of the system centres on the ESP32 development board, which provides dual-core processing, built-in Wi-Fi, sufficient GPIO for all sensor connections, and analog inputs for the moisture and LDR sensors. The DHT11 connects on a single digital pin using its one-wire protocol. Pump actuation runs through a MOSFET driver module that takes a 3.3V logic signal from the ESP32 and switches the 5V supply to the pump independently, avoiding any risk of back-EMF damage to the microcontroller. An external DC adapter powers the pump circuit.

Firmware is written in Embedded C++ using the Arduino framework, compiled and uploaded through Arduino IDE. The Firebase ESP32 library manages authentication and JSON serialization for all cloud writes. On the software side, the AI backend uses Python with TensorFlow 2.x, scikit-learn for preprocessing, and Flask as the HTTP layer. Google Cloud Run hosts this backend as a containerized service with auto-scaling so it handles bursts of requests without manual intervention. The Streamlit dashboard and the Flutter mobile application both consume data exclusively through Firebase, keeping the architecture clean and avoiding any direct dependency on the Cloud Run endpoint from client devices. Version control throughout used Git and GitHub.

V. RESULTS AND DISCUSSION

Testing was conducted in a controlled indoor setup with one plant in a soil-filled container. Sensor readings were logged continuously over multiple days, the irrigation system was deliberately triggered by allowing the soil to dry below threshold, and the AI backend was queried on each sensing cycle. The key results across the major system components are summarized below.

A. Soil Moisture and Irrigation Performance

Across 30 independent test runs — each defined as a full dry-to-irrigated cycle where moisture started above threshold, was allowed to drop below 30%, triggered the pump, and

recovered — the pump activated correctly every single time, giving a 100% irrigation activation success rate. Response latency from the moment moisture crossed threshold to pump activation was consistently in the 5–10 second range, which includes the Firebase write, the Cloud Run read, and the ESP32 polling interval. Moisture readings recovered to safe levels within one to three watering cycles depending on initial dryness. Compared with a fixed schedule that had historically watered twice daily regardless of soil state, the threshold-based approach reduced estimated water use by roughly 22–25%.

B. AI Model Performance

The soil quality model achieved a test accuracy of 93.42%, which is notably higher than the training accuracy of 82.14%. This pattern is somewhat unusual and most likely reflects the fact that the test set, after stratified splitting, ended up slightly more representative of the class distributions the model learned best. In practice, the model correctly labelled medium soil — the most common condition during testing — in every evaluated cycle. Poor predictions appeared reliably when moisture was critically low combined with high temperature, and good predictions appeared when all parameters were within comfortable ranges. Recommendation quality tracked classification quality well: irrigation suggestions were contextually appropriate in all observed cases, crop suggestions (Maize recommended during moderate stable conditions, Millets during dry stress, Rice during high humidity) aligned with standard agronomic guidance, and soil nutrient suggestions correctly identified when organic supplementation would be beneficial.

C. Cloud and Communication Performance

Firebase synchronization delay across all monitored sessions was minimal — updates appeared on the dashboard and mobile application within roughly one second of the ESP32 writing new data. The Cloud Run backend maintained stable uptime throughout the testing period. Device health monitoring worked as designed: when the ESP32 was deliberately disconnected from power, both the web dashboard and mobile application updated the device status to OFFLINE within the timeout window, and the status restored to ONLINE immediately when the device reconnected. One notable early issue was that relay-based pump control triggered erratically at ESP32 startup because the relay coil received a brief floating signal before firmware initialization completed. Replacing the relay with a MOSFET driver eliminated this entirely.

D. Overall System Integration

The full data path — from sensor reading on the ESP32 to a prediction visible on the mobile app — worked reliably end-to-end throughout all testing sessions. The architecture kept each layer independent enough that failures in one component did not cascade: when Cloud Run was briefly unavailable during one session, the ESP32 continued reading sensors,

Firestore continued storing data, and the dashboard continued displaying the last known values. Only the AI recommendations stalled temporarily, which is an acceptable degradation. Figure 6 in the project report shows the complete system integration schematic with all components and their data flows annotated.

VI. CONCLUSION

HillSense AI demonstrates that it is possible to build a genuinely useful smart agriculture system at a cost and complexity level that does not require specialized infrastructure or large-team engineering effort. The combination of a continuously sensing ESP32 node, Firestore as a real-time data bus, a TensorFlow model deployed on Cloud Run, and client interfaces on both web and mobile produces a system where the farmer always knows what their soil is doing and the soil is always being irrigated at the right moment rather than on a fixed schedule.

The 93.42% test accuracy on soil quality classification and the 100% irrigation activation rate across all test runs give us reasonable confidence that the core logic is sound. The 22–25% reduction in estimated water usage is not a trivial number in an agricultural context where water costs money and availability is becoming a more serious constraint every year. The MOSFET driver change that resolved early relay instability is worth highlighting because it is exactly the kind of practical engineering decision that rarely appears in published work but matters enormously in actual deployment.

Several directions for future work are clear. Adding a proper NPK sensor and a pH probe would make the soil nutrient recommendations quantitative rather than categorical. Integrating a weather forecast API would allow the system to anticipate rainfall and skip irrigation pre-emptively. Replacing the DHT11 with a more accurate sensor — it has a reputation for drift at humidity extremes — would improve recommendation reliability in humid conditions. On the software side, an LSTM-based model trained on time-series sensor histories rather than individual snapshots could capture diurnal patterns that the current classification model misses. Finally, a field deployment with multiple nodes and a mesh network would test whether the architecture scales to the kind of multi-farm scenario that precision agriculture ultimately requires.

VII. REFERENCES

- [1] A. Bahga and V. Madiseti, *Internet of Things: A Hands-On Approach*. Universities Press, 2015.
- [2] K. G. Liakos et al., “Machine Learning in Agriculture: A Review,” *Sensors*, vol. 18, no. 8, p. 2674, 2018.
- [3] A. Kamlaris and F. X. Prenafeta-Boldú, “Deep Learning in Agriculture: A Survey,” *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, 2018.
- [4] S. Wolfert et al., “Big Data in Smart Farming: A Review,” *Agricultural Systems*, vol. 153, pp. 69–80, 2017.
- [5] Google Developers, *Firestore Official Documentation*, 2024. [Online]. Available: <https://firebase.google.com/docs>

- [6] O. Elijah et al., "IoT and Data Analytics in Agriculture," IEEE Internet of Things Journal, vol. 5, no. 5, pp. 3758–3773, 2018.
- [7] A. Khanna and S. Kaur, "Evolution of IoT and Its Impact in Precision Agriculture," Computers and Electronics in Agriculture, vol. 157, pp. 218–231, 2019.
- [8] A. Géron, Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow. O'Reilly Media, 2019.
- [9] F. Chollet, Deep Learning with Python. Manning Publications, 2021.
- [10] TensorFlow Developers, TensorFlow Official Documentation, 2024. [Online]. Available: https://www.tensorflow.org/api_docs
- [11] Flutter Developers, Flutter Official Documentation, 2024. [Online]. Available: <https://docs.flutter.dev>
- [12] Espressif Systems, ESP32 Technical Reference Manual, 2024. [Online]. Available: <https://docs.espressif.com>
- [13] Streamlit Inc., Streamlit Official Documentation, 2024. [Online]. Available: <https://docs.streamlit.io>
- [14] Google Cloud, Cloud Run Documentation, 2024. [Online]. Available: <https://cloud.google.com/run/docs>
- [15] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.
- [16] W. McKinney, Python for Data Analysis. O'Reilly Media, 2022.
- [17] Arduino, Arduino Official Documentation, 2024. [Online]. Available: <https://docs.arduino.cc>
- [18] Plotly Technologies Inc., Plotly Python Documentation, 2024. [Online]. Available: <https://plotly.com/python>
- [19] S. Wolfert, L. Ge, C. Verdouw, and M. J. Bogaardt, "Big Data in Smart Farming," Agricultural Systems, vol. 153, pp. 69–80, 2017.
- [20] O. Elijah, T. A. Rahman, I. Orikumhi, C. Y. Leow, and M. N. Hindia, "An Overview of IoT and Data Analytics in Agriculture," IEEE Internet of Things Journal, vol. 5, no. 5, pp. 3758–3773, 2018.